Posterior Predictive Design for Phase I Clinical Trials

Appendix

A Predictive Bayes factor

The predictive Bayes factor is a generic and coherent[1] Bayesian hypothesis testing method (Zhou 2011), defined below. Denote $H_k(\mathbf{D})$ the posterior predictive model under hypotheses H_k, $k = 1, 2$, with the parameter updated in the presence of observed data $\mathbf{D} = \{y_1, y_2, \cdots, y_n\}$. Under mild regularity conditions, a comparison between $H_1(\mathbf{D})$ and $H_2(\mathbf{D})$ can be conducted using

$$\mathrm{PrBF}_{1,2} = \frac{\prod_i P(y_i|H_1(\mathbf{D}))}{\prod_i P(y_i|H_2(\mathbf{D}))} \frac{\exp(\hat{b}_1)}{\exp(\hat{b}_2)},$$

where $\hat{b}_k = -tr\{J_{n,k}^{-1}(\hat{\theta}^k) I_{n,k}(\hat{\theta}^k)\}$ is an asymptotically unbiased estimator for hypothesis k that corrects the estimation bias in the empirical log posterior predictive distribution, with the model-specific density function $g(y|\theta)$, prior distribution $\pi(\theta)$, posterior mode $\hat{\theta}$ and

$$\begin{aligned} J_n(\theta) &= -\frac{1}{n}\sum_{i=1}^{n}\left(\frac{\partial^2 \log\{g(y_i|\theta)\pi^{\frac{1}{n}}(\theta)\}}{\partial\theta\partial\theta'}\right), \\ I_n(\theta) &= \frac{1}{n}\sum_{i=1}^{n}\left(\frac{\partial \log\{g(y_i|\theta)\pi^{\frac{1}{n}}(\theta)\}}{\partial\theta} \frac{\partial \log\{g(y_i|\theta)\pi^{\frac{1}{n}}(\theta)\}}{\partial\theta'}\right). \end{aligned}$$

The predictive Bayes factor is an empirical estimator of $P(H_1(\mathbf{D}))/P(H_2(\mathbf{D}))$, the ratio of the posterior weights of posterior predictive models, given the equal (non-informative) prior model weights. In general, the predictive Bayes factors bear similarity to Bayes factors in interpretation by quantifying the "weight of evidence" in favor of one hypothesis against another. However, the evidence is assessed using the posterior predictive distribution rather than the prior predictive distribution, which can help to significantly reduce the sensitivity to prior variations and avoid the degeneration of the integrated likelihood underlying Lindley's paradox. Different from the posterior Bayes factors (Aitkin 1991) that

[1] In the context of model selection (e.g., in Johnson 2005), rather than in the context of clinical trial dose finding (Cheung 2005).

may improperly use data twice for predictive inference, predictive Bayes factors correct the asymptotic over-estimation error of the in-sample log posterior predictive distribution. More importantly, when the total sample size is small, such as in the setting of dose-finding trial designs, the predictive Bayes factors can also provide reasonable empirical performance without further correction. Given these merits, we adopt the predictive Bayes factor in hypothesis testing for the PoP design.

B Proof of Theorem 2

Proof. Denote $L_{n_j} = \frac{\arg\min_{y_j} \mathrm{PrBF}_{0,1}(y_j,n_j)\geq C}{n_j}$ and $U_{n_j} = \frac{\arg\max_{y_j} \mathrm{PrBF}_{0,1}(y_j,n_j)\geq C}{n_j}$ the lower and upper boundaries of the PoP design with n_j patients treated at dose level j, for $\forall j \in J$. Using the same notation as the main text by suppressing the subscript j, we first prove the following lemmas:

Lemma 1. *The predictive Bayes factor (3) converges to e as $n \to \infty$ if H_0 in (1) is true. Otherwise, if H_1 is true, it converges to 0.*

Proof. As $n \to \infty$,

(i) if H_0 is true, y approximates $n\phi_0$. Then, $\lim_{n\to\infty} \mathrm{PrBF}_{0,1} = e$

(ii) if H_1 is true, denote the maximum likelihood estimator (MLE) of ϕ be $\hat{\phi} = y/n$ such that

$$\lim_{n\to\infty} \mathrm{PrBF}_{0,1} = \lim_{n\to\infty} e \left\{ \left(\frac{\phi_0}{\hat{\phi}}\right)^{\hat{\phi}} \left(\frac{1-\phi_0}{1-\hat{\phi}}\right)^{(1-\hat{\phi})} \right\}^n$$

Note that $\left(\frac{\phi_0}{\hat{\phi}}\right)^{\hat{\phi}} \left(\frac{1-\phi_0}{1-\hat{\phi}}\right)^{(1-\hat{\phi})} \leq 1$ with equality holds if and only if $\hat{\phi} = \phi_0$. Thus, it does not hold under H_1. Then, $\lim_{n\to\infty} \mathrm{PrBF}_{0,1} = 0$.

□

Lemma 2. *Under the condition* $C < e$, $\lim_{n\to\infty} L_n = \phi$ *and* $\lim_{n\to\infty} U_n = \phi$.

Proof. Without loss of generality, we prove $\lim_{n\to\infty} L_n = \phi$. The proof for the upper bound is similar.

Since $\lim_{n\to\infty} \mathrm{PrBF}_{0,1}(n\phi, n) = e > C$, for $\epsilon = e - C > 0$, there exists an N, such that for $\forall n > N$, $\mathrm{PrBF}_{0,1}(n\phi, n) > e - \epsilon$. Then, as n is large enough, at least $\hat{y} = n\phi \in \{y : \mathrm{PrBF}_{0,1}(y, n) \geq C\}$. Thus, L_n exists.

Then, we prove $\lim_{n\to\infty} L_n = \phi$ by contradiction. If $\lim_{n\to\infty} L_n = \theta < \phi$, $\mathrm{PrBF}_{0,1}(n\theta, n) \geq C$. However, according to Lemma 1, $\lim_{n\to\infty} \mathrm{PrBF}_{0,1}(n\theta, n) = 0 < C$. Thus, $\lim_{n\to\infty} L_n = \phi$.

□

Let $\mathcal{J}$ denote the set of indices for dose levels that appear an infinite number of times in the sequential allocation such that $\mathcal{J} = \{j : n_j \to \infty \text{ as } n \to \infty\}$. Denote j^* the MTD as $j^* \equiv \mathrm{argmin}_{1\leq j\leq J}|\pi_j - \phi|$.

If $\pi_{j^*} = \phi$, we prove $P(j^* \in \mathcal{J}) = 1$ by contradiction. Assume there is a dose $k > j^*$ for which only dose levels above k are visited an infinite number of times. We know that $\pi_k > \phi$. According to Lemma 2, for a large enough n, we have $PrBF_{0,1}(y_k) < C$ and $\hat{\pi}_k > \phi$. This indicates that the next cohort of patients will be assigned to lower level $k-1$. Then, dose level $k-1$ should appear an infinite number of times as well, reaching a contradiction. Thus, there is no $k > j^*$ such that $k \in \mathcal{J}$. Similarly, there is no $k < j^*$ such that $k \in \mathcal{J}$. Thus, we have $j^* \in \mathcal{J}$ and for n large enough, $P(PrBF_{0,1}(y_{j^*}) \geq C \text{ and } \hat{\pi}_{j^*} = \phi) = 1$. Upon visit j^*, it will be repeatedly allocated with probability 1. Thus, the design will converge almost surely.

If no dose level has the toxicity rate ϕ, we prove that $P(j^*, j^*+1 \in \mathcal{J}) = 1$ if $\pi_{j^*} < \phi$ and $\pi_{j^*+1} > \phi$. The proof is similar when $\pi_{j^*-1} < \phi$ and $\pi_{j^*} > \phi$. Similar to the case

when $\pi_{j^*} = \phi$, we can prove that there is no $k > j^* + 1$ such that $k \in \mathcal{J}$ and there is also no $k < j^*$ such that $k \in \mathcal{J}$. For a large enough n, once visit j^*, we have $PrBF_{0,1}(y_{j^*}) < C$ and $\hat{\pi}_{j^*} < \phi$. Thus, the next cohort of patients will be assigned to dose j^*+1. Similarly, once visit $j^* + 1$, the next cohort of patients will be assigned to dose j^*. Thus, the design would eventually oscillate between two dose levels at which the associated toxicity rates straddle ϕ.

If all π_j's are above or below ϕ, similarly, the dose levels straddle ϕ, including j^*, will be repeatedly allocated with probability 1. According to the Lemma 1 of Oron et al. (2011), the design will converge almost surely.

□

C Proof of Theorem 3

Proof. Without loss of generality, we prove $|L_n - \phi| \leq kn^{-1/2}$ as $n \to \infty$, where $k = \sqrt{2\phi(1-\phi)(1-\log C)}$. The proof for $|U_n - \phi| \leq kn^{-1/2}$ is similar.

Let $f(x) = n\left[x \log \frac{\phi}{x} + (1-x) \log \frac{1-\phi}{1-x}\right]$, such that $\text{PrBF}_{0,1} = \exp\left(f\left(\frac{y}{n}\right) + 1\right)$. Then, we have

$$f(L_n) \geq \log C - 1,$$

$$f(\phi) = f'(\phi) = 0,$$

$$f''(\phi) = -\frac{n}{\phi(1-\phi)}.$$

Based on the Taylor theorem, there exists a function $h_2 : \mathcal{R} \to \mathcal{R}$, such that

$$\begin{aligned} f(x) &= f(\phi) + f'(\phi)(x-\phi) + f''(\phi)(x-\phi)^2 + h_2(x)(x-\phi)^2 \\ &= \left(-\frac{n}{2}\frac{1}{\phi(1-\phi)} + h_2(x)(x-\phi)^2\right) \end{aligned}$$

and $\lim_{x\to\phi} h_2(x) = 0$. Then, according to Lemma 2,

$$f(L_n) = \left(-\frac{n}{2}\frac{1}{\phi(1-\phi)} + h_2(L_n)(L_n-\phi)^2 \right) \geq \log C - 1$$

. Thus,

$$(L_n - \phi)^2 \leq \frac{2\phi(1-\phi)(1-\log C)}{n}$$

. it gives

$$|L_n - \phi| \leq \sqrt{2\phi(1-\phi)(1-\log C)}n^{-1/2}$$

□

D Proof of Theorem 4

Proof. According to the decision rules, the probability of dose escalation of the PoP design when the observed DLT rate is $\pi_j < \phi$ is given by $P(\text{dose escalation}|\hat{\pi}_j > \phi) = P(PrBF_{0,1} < C \ \& \ \hat{\pi}_j < \phi \,|\, \hat{\pi}_j > \phi) = 0$. Thus, the PoP design will not escalate the dose when the observed DLT rate is higher than the target.

Similarly, $P(\text{dose de-escalation}|\hat{\pi}_j < \phi) = P(PrBF_{0,1} < C \ \& \ \hat{\pi}_j > \phi \,|\, \hat{\pi}_j < \phi) = 0$. Thus, the PoP design never de-escalates the dose when the observed DLT rate is lower than the target toxicity rate. So, the PoP design is long-term memory coherent.

□

E Details of implementing CRM

For the CRM method, we obtained the skeleton using the *getprior* function from the R package *dfcrm*, employing the indifference-interval based approach of Lee and Cheung (2009). We set the halfwidth of the indifference intervals to 0.05, used the middle dose level (i.e., dose level 2 for K=4 and dose level 3 for K=6) as the prior guess of MTD, and employed a one-parameter logistic model as the working model. Specifically,

- when $\phi = 0.2$, the skeleton is $(0.112, 0.200, 0.311, 0.429)$ for $K = 4$,
  and is $(0.055, 0.112, 0.200, 0.311, 0.429, 0.539)$ for $K = 6$;

- when $\phi = 0.25$, the skeleton is $(0.158, 0.250, 0.355, 0.462)$ for $K = 4$,
  and is $(0.089, 0.158, 0.250, 0.355, 0.462, 0.558)$ for $K = 6$;

- when $\phi = 0.3$, the skeleton is $(0.205, 0.300, 0.402, 0.500)$ for $K = 4$,
  and is $(0.126, 0.205, 0.300, 0.402, 0.500, 0.587)$ for $K = 6$.

F Design performance in additional scenarios

Table A1 shows the four prespecified scenarios for oncology trials used in Liu and Yuan (2015) but with a different target toxicity rate. These scenarios were originally proposed for $\phi = 0.25$, but here we examined when $\phi = 0.20$ to demonstrate the operating characteristics of the PoP, BOIN, Keyboard, and CRM designs when the target toxicity rate may fall in between two adjacent dose levels. MTD is considered the dose level whose DLT rate is the closest to the target rate. As expected, the parametric CRM design performs the best if the model assumptions are satisfied and the dose-toxicity skeleton closely resembles the truth. Otherwise, if the assumptions are violated, the CRM design will not perform as well, suggesting the importance of conducting comprehensive assessments across multiple scenarios (e.g. in Table 2). Simulation studies find that PoP design has excellent performance in all settings and outperforms other interval designs.

Table A1: Performance metrics of the PoP, BOIN, Keyboard, and CRM designs under four prespecified dose-toxicity scenarios for cohort sizes of 1. The target toxicity rate is 0.20. $N = 36$. Results from isolated scenarios should be interpreted with caution, as performance comparisons across different design classes can substantially vary by scenario.

Design		Target ϕ=0.20							
		Dose level						Not Choosing Any Dose	Risk of Over-dosing
		1	2	3	4	5	6		
Scenario 1	Pr(toxicity)	0.25	0.35	0.5	0.6	0.7	0.8		
PoP	Selection(%)	**65.0**	9.7	0.4	0.0	0.0	0.0	24.9	7.8
	#Patients	**22.8**	6.4	1.7	0.5	0.2	0.0		
BOIN	Selection(%)	49.5	9.4	0.4	0.0	0.0	0.0	40.7	9.7
	#Patients	17.3	6.4	2.1	0.7	0.2	0.1		
Keyboard	Selection(%)	49.7	8.8	0.4	0.0	0.0	0.0	41.1	8.6
	#Patients	17.9	6.0	1.9	0.6	0.2	0.0		
CRM	Selection(%)	48.9	8.5	0.3	0.0	0.0	0.0	42.3	7.3
	#Patients	18.3	5.7	1.6	0.5	0.1	0.0		
Scenario 2	Pr(toxicity)	0.1	0.25	0.4	0.6	0.7	0.8		
PoP	Selection(%)	37.6	55.7	5.4	0.1	0.0	0.0	1.3	2.6
	#Patients	13.8	15.2	4.8	1.0	0.3	0.1		
BOIN	Selection(%)	42.2	48.8	5.2	0.2	0.0	0.0	3.6	3.5
	#Patients	14.2	13.9	5.0	1.4	0.4	0.1		
Keyboard	Selection(%)	43.1	48.5	4.6	0.1	0.0	0.0	3.7	3.1
	#Patients	14.9	13.8	4.6	1.2	0.3	0.1		
CRM	Selection(%)	33.5	**57.2**	5.6	0.0	0.0	0.0	3.6	2.4
	#Patients	13.5	**15.7**	4.6	0.9	0.2	0.1		
Scenario 3	Pr(toxicity)	0.05	0.1	0.25	0.32	0.5	0.6		
PoP	Selection(%)	2.4	36.5	47.1	13.1	0.6	0.0	0.3	8.3
	#Patients	3.6	11.5	12.3	5.9	1.7	0.4		
BOIN	Selection(%)	4.0	39.4	41.0	14.0	0.9	0.1	0.6	9.8
	#Patients	3.6	11.8	11.1	6.4	2.1	0.7		
Keyboard	Selection(%)	4.4	40.2	41.1	13.0	0.6	0.0	0.6	9.3
	#Patients	3.9	12.3	11.1	6.1	1.9	0.6		
CRM	Selection(%)	1.0	32.2	**53.9**	11.9	0.4	0.0	0.6	8.8
	#Patients	3.1	11.1	**14.0**	5.8	1.4	0.4		
Scenario 4	Pr(toxicity)	0.01	0.02	0.03	0.04	0.05	0.25		
PoP	Selection(%)	0.0	0.1	0.3	1.1	31.2	67.2	0.0	0.0
	#Patients	1.2	1.4	1.7	2.4	11.2	17.3		
BOIN	Selection(%)	0.1	0.2	0.4	1.0	40.8	57.4	0.0	0.0
	#Patients	1.2	1.3	1.5	2.1	13.2	16.7		
Keyboard	Selection(%)	0.1	0.4	0.5	1.5	41.3	56.2	0.0	0.0
	#Patients	1.2	1.4	1.6	2.2	13.4	16.3		
CRM	Selection(%)	0.0	0.0	0.2	1.3	26.9	**71.6**	0.0	0.0
	#Patients	1.1	1.2	1.6	2.5	10.4	**19.2**		

G Dose selection

Here we investigated the distributed dose selection performance of the PoP, BOIN, Keyboard, and CRM designs. In addition to the chance of correct selection of MTD (% Correct Selection), we calculated the percentage of selecting a dose above MTD (% Overdose Selection), under MTD (% Underdose Selection), and claiming no MTD if the lowest dose is considered overly toxic even though MTD exists (% No Selection).

Figure A1 presents the simulation results. Among all the designs, the PoP design exhibited the highest % correct selection and the lowest % no selection, indicating its efficiency and robustness in MTD identification. The CRM design demonstrated the highest % overdosing selection and lowest % underdosing selection, suggesting a tendency to select a higher dose as MTD. In contrast, the Keyboard design was more likely to select a dose level lower than the true MTD, resulting in not only the lowest % Overdose Selection but also the lowest % Correct Selection. Overall, this difference among the designs highlights their risk preference and OC in MTD selection.

Figure A1: Barplots of the percentage of overdose, correct, underdose, and no selection for the PoP, BOIN, Keyboard, and CRM designs over 10,000 simulated scenarios.

H Simulation of consistency

The consistency property was proved in Appendix D. A simulation study with $K = 3$ and $phi = 0.25$ various sample size n was examined to empirically demonstrate the property. As Figure A2 illustrates, the PCS increases to 1 as the sample size increases, indicating the consistency of the PoP design.

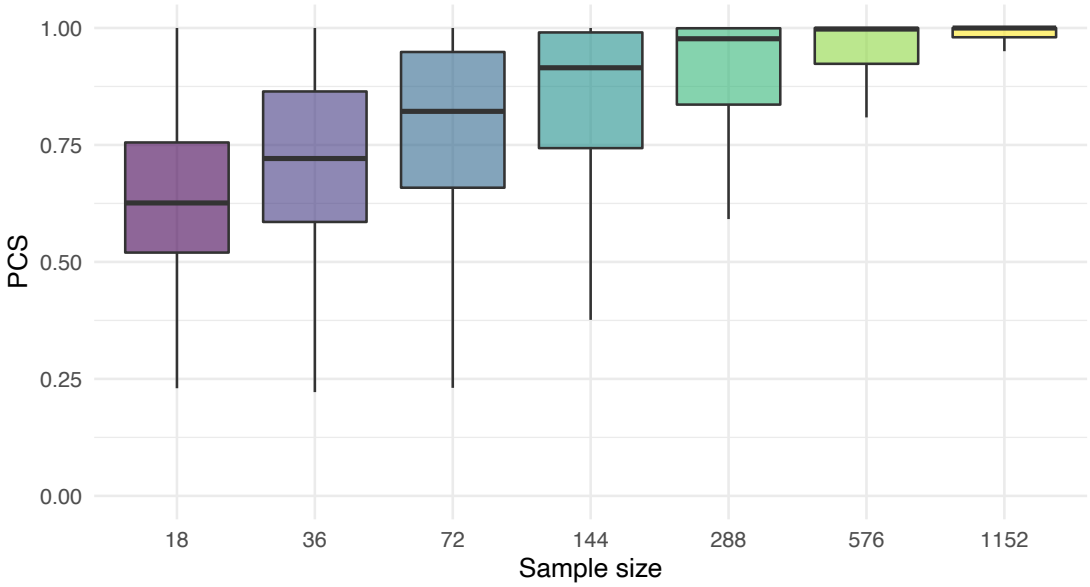

Figure A2: Boxplots of the percentage of correct selection (PCS) for the PoP design for different sample sizes over 20,000 simulated scenarios. Dose exclusion and elimination rules are deactivated for illustration purposes.

I Sensitivity analysis of PoP design for loss functions

In this sensitivity analysis, we investigate the performance of the PoP design with respect to various sets of loss functions. We restrict the sensitivity analysis to the setting of $K = 3$ doses and $n = 18$ patients, so as not to lose the main ideas. Under the conditions mentioned in Section 4, we examine $b_1 = 0.2$ or 0.25, $b_2 = 0.6$ or $2/3$, and $b_3 = 0.15$ or $1/6$.

Table A2 shows that the PoP designs generate universally higher PCS and PCA than the BOIN and Keyboard designs. Different choices of loss functions punish the dose transition to various extents so that the risk of overdosing at 70% varies.

Table A2: Percentage of correct selection (PCS, %), percentage of correct allocation (PCA, %) and overdose control for different loss functions.

K=3,N=18				ϕ=0.15			ϕ=0.2			ϕ=0.25		
	b1	b2	b3	PCS	PCA	Risk of OD	PCS	PCA	Risk of OD	PCS	PCA	Risk of OD
BOIN				56.6	48.8	11	56.8	48.0	10.1	57.3	48.6	9.7
Keyboard				56.0	49.2	11.2	56.6	48.4	10	56.0	48.0	9.5
PoP	0.2	0.6	0.2	59.8	52.2	9.9	61.0	52.0	10.4	61.3	51.7	11.1
	0.2	0.6	1/6	59.7	52.3	10.5	60.0	51.4	10.4	61.1	51.5	10.0
	0.2	2/3	0.2	60.0	51.5	8.1	61.3	51.1	9.2	61.5	50.9	9.7
	0.2	2/3	1/6	60.4	52.0	9.2	60.2	50.6	9.1	61.5	51.4	10.6
	0.25	0.6	0.2	60.1	53.1	13.1	59.2	51.5	11.1	60.8	52.3	14.2
	0.25	0.6	1/6	60.1	53.1	13.1	59.6	51.9	13.6	60.5	52.1	12.7
	0.25	2/3	0.2	59.7	52.4	10.5	60.1	51.7	11.8	60.7	51.8	10.7
	0.25	2/3	1/6	60.4	53.0	11.8	59.7	51.7	11.8	61.1	52.2	13.3

K=3,N=18				ϕ=0.3			ϕ=0.35			ϕ=0.4		
	b1	b2	b3	PCS	PCA	Risk of OD	PCS	PCA	Risk of OD	PCS	PCA	Risk of OD
BOIN				57.3	48.6	10.2	58.4	49.1	10.5	56.8	47.9	8.2
Keyboard				55.1	47.4	10.1	54.7	46.2	8.7	56.8	47.2	9.6
PoP	0.2	0.6	0.2	60.9	51.3	10.9	61.7	51.9	12.0	61.8	51.6	10.2
	0.2	0.6	1/6	60.4	51.2	11.4	61.6	51.7	11.9	61.6	51.8	10.8
	0.2	2/3	0.2	60.9	50.2	9.2	61.7	50.6	8.7	61.7	50.2	9.2
	0.2	2/3	1/6	60.4	50.1	9.2	61.6	50.4	8.6	61.6	50.6	10.5
	0.25	0.6	0.2	60.3	51.5	12.2	61.1	52.3	12.0	61.4	52.2	12.0
	0.25	0.6	1/6	59.8	51.2	12.9	61.0	52.0	11.9	61.1	52.1	11.2
	0.25	2/3	0.2	60.7	51.5	12.1	61.5	52.0	12.5	61.6	51.9	11.2
	0.25	2/3	1/6	60.2	51.2	12.1	61.4	51.9	11.4	61.5	51.9	11.2

Abbreviation: OD: over-dosing.

References

Aitkin, M. (1991), 'Posterior bayes factors', *Journal of the Royal Statistical Society: Series B (Methodological)* **53**(1), 111–128.

Cheung, Y. K. (2005), 'Coherence principles in dose-finding studies', *Biometrika* **92**(4), 863–873.

Johnson, V. E. (2005), 'Bayes factors based on test statistics', *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* **67**(5), 689–701.

Lee, S. M. & Cheung, Y. K. (2009), 'Model calibration in the continual reassessment method', *Clinical Trials* **6**(3), 227–238.

Liu, S. & Yuan, Y. (2015), 'Bayesian optimal interval designs for phase i clinical trials', *Journal of the Royal Statistical Society. Series C: Applied Statistics* **64**, 507–523.

Oron, A. P., Azriel, D. & Hoff, P. D. (2011), 'Dose-finding designs: The role of convergence properties', *International Journal of Biostatistics* **7**.

Zhou, S. (2011), Bayesian Model Selection in terms of Kullback-Leibler discrepancy, PhD thesis, Columbia University.